Bayesian variable selection for latent class analysis using a collapsed Gibbs sampler

Arthur White*, Jason Wyse†, and Thomas Brendan Murphy‡§

arXiv:1402.6928v1 [stat.CO] 27 Feb 2014

Abstract

Latent class analysis is used to perform model based clustering for multivariate categorical responses. Selection of the variables most relevant for clustering is an important task which can affect the quality of clustering considerably. This work considers a Bayesian approach for selecting the number of clusters and the best clustering variables. The main idea is to reformulate the problem of group and variable selection as a probabilistically driven search over a large discrete space using Markov chain Monte Carlo (MCMC) methods. This approach results in estimates of degree of relevance of each variable for clustering along with posterior probability for the number of clusters. Bayes factors can then be easily calculated, and a suitable model chosen in a principled manner. Both selection tasks are carried out simultaneously using an MCMC approach based on a collapsed Gibbs sampling method, whereby several model parameters are integrated from the model, substantially improving computational performance. Approaches for estimating posterior marginal probabilities of class membership, variable inclusion and number of groups are proposed, and post-hoc procedures for parameter and uncertainty estimation are outlined. The approach is tested on simulated and real data.

1 Introduction

Latent class analysis (LCA) models (Goodman, 1974) are used to discover affinities and groupings in multivariate categorical response data. An example of data for which a LCA model may be appropriate would be records for N items on M different variables where each of the variables observed is categorical, more precisely, binary or nominal. There may be a varying number of categories over the M different variables. In using a LCA model one expects the items to segment into contiguous groups within which items are similar in

*School of Computer Science and Statistics, O'Reilly Institute, Trinity College, Dublin 2, Ireland. Work performed while in School of Mathematical Sciences, Complex & Adaptive Systems Laboratory and Insight Research Centre, University College Dublin, Dublin 4, Ireland.

†School of Computer Science and Statistics, O'Reilly Institute, Trinity College, Dublin 2, Ireland.

‡School of Mathematical Sciences, Complex & Adaptive Systems Laboratory and Insight Research Centre, University College Dublin, Dublin 4, Ireland.

§This work is supported by Science Foundation Ireland under the Clique Strategic Research Cluster (SFI/08/SRC/I1407) and the Insight Research Centre (SFI/12/RC/2289).

nature. These groupings are captured by allowing the variables' multinomial probabilities to vary by group.

LCA models have been widely studied and applied (Aitkin et al, 1981; Garrett and Zeger, 2000; Walsh, 2006). They can be viewed as a special case of model-based clustering (McLachlan and Peel, 2002; Fraley and Raftery, 2007), in which each item is assumed to arise from one of a number of groups, each group having its own data probability distribution. As with other applications of clustering, two of the main difficulties when formulating a LCA model are identifying a suitable number of different groups or clusters in the data, and choosing the variables which are most informative for groupings in the data.

Dean and Raftery (2010) have demonstrated that both of these aspects of model selection can have a considerable effect on the resulting clustering of the data, and that these questions are important in the formulation and application of LCA models. Their approach learns from the data using a headlong algorithm search to choose the optimal clustering variables and number of groups, with steps in the algorithm determined using the Bayesian Information criterion (BIC) (Schwarz, 1978; Kass and Raftery, 1995).

In this paper, a Bayesian approach for the analysis of LCA models is proposed. It is shown how simple marginalization of the parameters in a LCA model leads to a form of the model for which Markov chain Monte Carlo (MCMC) sampling algorithms can be used to quantify precisely the uncertainty in the number of groups in the data, as well as which variables give the best clustering. This is similar to work carried out by Nobile and Fearnside (2007) in the analysis of Gaussian finite mixtures, Wyse and Friel (2012) in a block clustering application, McDaid et al (2013) in the clustering of social network data, and Tadesse et al (2005) in the context of clustering variable selection for Gaussian distributed data.

The advantage of such an approach is that it leads to estimates of the posterior probability of a variable being useful for clustering the data as will be demonstrated. This gives more information than would be available from the headlong search algorithm in Dean and Raftery (2010) which gives a point estimate of the variables which optimally determine the clustering, but does not include any quantification of the uncertainty around these particular choices.

The rest of this paper is organised as follows. Section 2 outlines model specification for LCA in detail and gives a brief overview of other approaches for analysis that have been proposed in the literature. The four datasets which the method is applied to are then introduced and described. Section 3 presents the adopted marginalization approach and gives details of the MCMC algorithm used to estimate the models. Section 4 discusses estimation of posterior quantities from the model, including parameter estimation, how model comparison may be performed by directly calculation of Bayes factors, as well as a description of how label switching is accounted for. Section 5 applies the sampling algorithm to simulated and real data, concluding with a discussion in Section 6. Some additional results are included in supplementary material available online.

2 The classic LCA model

Denote the data by an $N \times M$ matrix $\mathbf{X}$ where each row is a record of responses for M categorical variables for one item. Row n of $\mathbf{X}$ is $\mathbf{X}_n = (X_{n1}, X_{n2}, \ldots, X_{nM})$. The entry X_{nm} takes a value from one of the categories $\{1, 2, \ldots, C_m\}$ for variable m. The LCA model

assumes that the $\mathbf{X}_n$ arise independently from a finite mixture model with G components,

$$p(\mathbf{X}_n|\boldsymbol{\tau},\boldsymbol{\theta}) = \sum_{g=1}^{G} \tau_g p(\mathbf{X}_n|\boldsymbol{\theta}_g).$$

The τ_g are mixture weights, with $\sum_{g=1}^{G} \tau_g = 1$ and $\tau_g > 0$, for all g. Each component has its own set of parameters $\boldsymbol{\theta}_g$ which embody the differences between groups; this holds the multinomial probabilities for all variables for class g. The parameter θ_{gmc} corresponds to the probability that an item takes category c for variable m within class g. For the component data likelihood, a *local independence* assumption is made. This assumes that, conditional on an observation's group membership, all variables are independent of each other, so that

$$p(\mathbf{X}_n|\boldsymbol{\theta}_g) = \prod_{m=1}^{M} \prod_{c=1}^{C_d} \theta_{gmc}^{\mathrm{I}(X_{nm}=c)}.$$

It is convenient to work with the completed data, that is, data augmented with class labels for each item. Denote these by $\mathbf{Z}$, where

$$Z_{ng} = \begin{cases} 1 & \text{if observation } n \text{ belongs to group } g \\ 0 & \text{otherwise.} \end{cases}$$

The completed data likelihood for an observation may then be written

$$p(\mathbf{X}_n\mathbf{Z}_n|\boldsymbol{\theta},\boldsymbol{\tau},G) = \prod_{g=1}^{G} \left\{ \tau_g \prod_{m=1}^{M} \prod_{c=1}^{C_d} \theta_{gmc}^{\mathrm{I}(X_{nm}=c)} \right\}^{Z_{ng}}.$$

2.1 Approaches for analysis

When the number of groups and variables is assumed fixed, several techniques for performing LCA are available. When attempting to identify the number of groups in the data, models are fitted over a range of groups, with the best fit often determined with the use of an information criterion.

In a frequentist paradigm, an expectation-maximisation (EM) algorithm (Dempster et al, 1977) can be employed. The BIC or Akaike's information criterion (AIC) (Akaike, 1973) can then be used to identify an optimal model. Goodness-of-fit statistics such as the likelihood ratio test (Goodman, 1974) can also be employed, but may prove difficult to apply to sparse data with a large number of variables (Aitkin et al, 1981).

In a Bayesian paradigm, a Gibbs sampling (Geman and Geman, 1984; Garrett and Zeger, 2000) technique can be used. In this setting, criterion such as the deviance information criterion (DIC) (Spiegelhalter et al, 2002), AIC or BIC Monte Carlo (AICM or BICM) (Raftery et al, 2007), have been proposed to aid model choice. One criticism of the DIC is that its specification is somewhat opaque; this can lead to different results depending on its interpretation when used in a mixture model setting (Celeux et al, 2006). Garrett and Zeger (2000) suggest graphical tools to aid model selection, such as weak identifiability, where a parameter's prior and posterior distributions are visually compared. If the difference

between these distributions is judged to be small, this then suggests that too large a number of groups has been fitted to the data.

Another approach for Bayesian model selection is to use reversible jump MCMC methods (Green, 1995). These have been extended to univariate and multivariate finite Gaussian mixture models (Richardson and Green, 1997; Dellaportas and Papageorgiou, 2006), but to our knowledge have never been applied to categorical data. Tadesse et al (2005) make use of this method to also perform variable selection, with some parameters integrated from the model. We base our approach on an alternative method, a fully collapsed sampler first proposed by Nobile and Fearnside (2007).

Technical issues can also occur when fitting LCA, such as underestimation of standard errors when using the EM algorithm (Bartholomew and Knott, 1999; Walsh, 2006), and label-switching for Gibbs sampling (Marin et al, 2005). The latter issues, and the methods used to avoid them, are discussed in Section 4.

2.2 Datasets to be analysed

We provide a description of the four datasets which are analysed in Section 5. The first example is simulated data, while two of the three real datasets have been previously analysed using the LCA methods described in Section 2.1. The last dataset we apply LCA methods to for the first time. In all four examples the data is binary, so that $C_m = 2,$ and $X_{nm} \in \{0, 1\}$ for all possible values of n and m. However, the software developed from the outlined method in Section 3 can just as easily be applied to categorical data where $C_m \geq 2$.

Dean and Raftery simulated dataset

We first test our approach on simulated data as specified by Dean and Raftery (2010). This is a binary two-class model of $N = 500$ observations with 4 informative variables (1-4) and 9 noise variables (5-13), with weights set to $\tau_1 = 0.6$ and $\tau_2 = 0.4$. The specified parameter values for $\boldsymbol{\theta}$ are given in Table 1.

Alzheimer dataset

This dataset of patient symptoms was recorded in the Mercer Institute of St. James' Hospital in Dublin, Ireland (Moran et al, 2004; Walsh, 2006). The data is a recording of the presence or absence of $M = 6$ symptoms displayed by $N = 240$ patients diagnosed with early onset Alzheimer's disease, and is available in the **R**(R Core Team, 2013) package **BayesLCA**(White and Murphy, 2013). Previous studies had difficulty in determining whether two or three groups are more suitable for the data, where fitting a three group model also created difficulties when performing inference (Walsh, 2006).

Teaching styles dataset

This dataset was conducted in an attempt to ascertain the different types of teaching style being employed in schools in the United Kingdom in the mid 1970s, and was discussed at length by Aitkin et al (1981) after an initial analysis by Bennet (1976). The dataset records which of $M = 39$ teaching methods are employed by the $N = 467$ schools. Computational

Table 1: Specified parameter values for $\boldsymbol{\theta}$ in the Dean and Raftery simulated dataset.

	Group 1	Group 2
1	0.6	0.2
2	0.8	0.5
3	0.7	0.4
4	0.6	0.9
5	0.5	0.5
6	0.4	0.4
7	0.3	0.3
8	0.2	0.2
9	0.9	0.9
10	0.6	0.6
11	0.7	0.7
12	0.8	0.8
13	0.1	0.1

limitations at the time meant that only a two or three group clustering of the data were seriously considered, with dimension reduction in a previous study performed using principal component analysis. Further, use of information based criterion such as the BIC had not been developed at that time, so that no automatic comparison of the clusterings could be made, with a decision on the optimal clustering being based on careful consideration of the different properties of the clusterings.

Physiotherapy dataset

We also use apply our methods to a physiotherapy dataset based on a survey recently conducted by Keith Smart in St. Vincent's University Hospital, Dublin (Smart et al, 2011). The aim of this study was to identify which symptoms distinguish between three types of back pain which patients suffer from. The dataset consists of $N = 425$ observations and $M = 36$ variables. While the different types of back pain were considered reasonably distinct, different subgroups of pain sufferers within these pain classes are also possible, which motivates an examination of the data in an unsupervised setting.

3 Bayesian latent class model and marginalization approach

In this section, a variable inclusion indicator variable is introduced, and a fully Bayesian specification for LCA is provided. A collapsed sampling scheme for inference is then described.

3.1 Clustering variable selection

The marginalization approach addresses exploring uncertainty in variable selection for clustering the items as well as uncertainty in the number of classes, G. To this end, introduce a binary vector $\boldsymbol{\nu} = (\nu_1, \ldots, \nu_M)$ for which ν_j takes the value 1 if j is included as a clustering variable. Let $\boldsymbol{\nu}_{\text{cl}}$ be the set of variables used in the clustering and $\boldsymbol{\nu}_{\text{n}} = \boldsymbol{\nu} \backslash \boldsymbol{\nu}_{\text{cl}}$. The variables which are not used for clustering should follow the same rule for all items, so that the data likelihood for these variables is

$$p_{\text{n}}(\mathbf{X}|\boldsymbol{\rho}, \boldsymbol{\nu}) = \prod_{n=1}^{N} \prod_{m \in \nu_{\text{n}}} \prod_{c=1}^{C_m} \rho_{mc}^{\text{I}(X_{nm}=c)}.$$

where ρ_{mc} is the probability of variable m having category c. Similarly, the completed data likelihood for the clustering variables is

$$\begin{aligned} p_{\text{cl}}(\mathbf{X}, \mathbf{Z}|\boldsymbol{\theta}, \boldsymbol{\nu}, \boldsymbol{\tau}, G) &= \prod_{n=1}^{N} \prod_{g=1}^{G} \left\{ \tau_g \prod_{m \in \nu_{\text{cl}}} \right. \\ &\times \left. \prod_{c=1}^{C_m} \theta_{gmc}^{\text{I}(X_{nm}=c)} \right\}^{Z_{ng}}. \end{aligned}$$

3.2 Prior assumptions and joint posteriors

Prior assumptions for all sets of multinomial probabilities as well as the component weights $\boldsymbol{\tau}$ are Dirichlet. For example, for clustering variable $m \in \boldsymbol{\nu}_{\text{cl}}$

$$p(\boldsymbol{\theta}_{gm}|\beta) = \frac{\Gamma\{C_m\beta\}}{\Gamma\{\beta\}^{C_m}} \prod_{c=1}^{C_m} \theta_{gmc}^{\beta-1}.$$

The hyperparameters β are chosen to be 1 for all g, m combinations. A Dirichlet prior is also assumed for those variables $m \in \nu_{\text{n}}$, where the hyperparameters β are again chosen to be 1 in all cases:

$$p(\boldsymbol{\rho}_m|\beta) = \frac{\Gamma\{C_m\beta\}}{\Gamma\{\beta\}^{C_m}} \prod_{c=1}^{C_m} \rho_{mc}^{\beta-1}.$$

For the component weights the prior assumed is

$$p(\boldsymbol{\tau}|\alpha, G) = \frac{\Gamma\{G\alpha\}}{\Gamma\{\alpha\}^{G}} \prod_{g=1}^{G} \tau_g^{\alpha-1}$$

where α is taken to be 0.5, for all possible values of g; this is the marginal Jeffreys for the multinomial model, and is partly chosen to discourage overfitting (Rousseau and Mengersen, 2011).

It is assumed that the prior probability of any variable being a clustering variable is π, giving a joint prior for $\boldsymbol{\nu}$ of

$$p(\boldsymbol{\nu}|\pi) = \prod_{m \in \nu_{\text{cl}}} \pi \prod_{m \in \nu_{\text{n}}} (1 - \pi).$$

A value of $\pi = 0.5$ is assumed for all the applications in this paper.

The complete data likelihood for all variables is

$$p_{\text{full}}(\mathbf{X}, \mathbf{Z}|\boldsymbol{\theta}, \boldsymbol{\rho}, \boldsymbol{\nu}, \boldsymbol{\tau}, G) = p_{\text{cl}}(\mathbf{X}, \mathbf{Z}|\boldsymbol{\theta}, \boldsymbol{\nu}, \boldsymbol{\tau}, G) p_{\text{n}}(\mathbf{X}|\boldsymbol{\rho}, \boldsymbol{\nu})$$

and the joint posterior for the unknowns (except the number of classes) is

$$\begin{aligned} p(\mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\rho}, \boldsymbol{\nu}, \boldsymbol{\tau}|\mathbf{X}, G, \alpha, \pi, \beta) &\propto p_{\text{full}}(\mathbf{X}, \mathbf{Z}|\boldsymbol{\theta}, \boldsymbol{\rho}, \boldsymbol{\nu}, \boldsymbol{\tau}, G) \\ &\times p(\boldsymbol{\tau}|\alpha, G) p(\boldsymbol{\nu}|\pi) \\ &\times \prod_{d \in \nu_{\text{n}}} p(\boldsymbol{\rho}_d|\beta) \\ &\times \prod_{g=1}^{G} \prod_{m \in \nu_{\text{cl}}} p(\boldsymbol{\theta}_{gm}|\beta). \end{aligned}$$

Under the model being considered there is also uncertainty in the number of classes. This is accounted for by taking a prior on G, which we assume to be Poisson(1), truncated at $G_{\max}$, the maximum number of classes considered feasible. The full posterior is then

$$p(G, \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\rho}, \boldsymbol{\nu}, \boldsymbol{\tau}|\mathbf{X}, \alpha, \pi, \beta) \propto p(\mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\rho}, \boldsymbol{\nu}, \boldsymbol{\tau}|\mathbf{X}, G, \alpha, \pi, \beta) p(G).$$

3.3 Marginalization approach

The marginalization approach proceeds by observing that all of the multinomial probabilities, as well as the component weights can be marginalized analytically out of the model by using the normalizing constant of the Dirichlet distribution. This leaves a joint distribution for G, $\boldsymbol{\nu}$ and $\mathbf{Z}$:

$$\begin{aligned} & p(G, \mathbf{Z}, \boldsymbol{\nu}|\mathbf{X}, \alpha, \pi, \beta) \\ \propto\ & p(G) \int p(\mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\rho}, \boldsymbol{\nu}, \boldsymbol{\tau}|\mathbf{X}, G, \alpha, \pi, \beta)\, d\boldsymbol{\theta}\, d\boldsymbol{\rho}\, d\boldsymbol{\tau}. \end{aligned}$$

Carrying out this marginalization gives

$$\begin{aligned} & p(G, \mathbf{Z}, \boldsymbol{\nu}|\mathbf{X}, \alpha, \pi, \beta) \\ \propto\ & p(G) p(\boldsymbol{\nu}) \frac{\Gamma\{G\alpha\}}{\Gamma\{\alpha\}^G} \frac{\prod_{g=1}^{G} \Gamma\{N_g + \alpha\}}{\Gamma\{N + G\alpha\}} \\ \times\ & \prod_{m \in \nu_{\text{n}}} \frac{\Gamma\{C_m\beta\}}{\Gamma\{\beta\}^{C_m}} \frac{\prod_{c=1}^{C_m} \Gamma\{N_{mc} + \beta\}}{\Gamma\{N + C_m\beta\}} \\ \times\ & \prod_{g=1}^{G} \prod_{m \in \nu_{\text{cl}}} \frac{\Gamma\{C_m\beta\}}{\Gamma\{\beta\}^{C_m}} \frac{\prod_{c=1}^{C_m} \Gamma\{N_{gmc} + \beta\}}{\Gamma\{N_g + C_m\beta\}} \end{aligned}$$

where N_g is the number of observations clustered to group g, N_{mc} is the number of times variable m takes category c, and N_{gmc} is the number of items in group g that have category c for variable m. Of note here is that this posterior makes sampling the number of groups G and the clustering variables $\boldsymbol{\nu}_{\text{cl}}$ possible using standard MCMC techniques as outlined in the following section.

3.4 Sampling algorithm

The sampling algorithm comprises three main operations. The first samples the class membership of observations, the second samples the number of classes and the third step samples probability of variable inclusion for clustering.

3.4.1 Class memberships

Class memberships are sampled using a Gibbs sampling step which exploits the full conditional distribution of the class label for observation n, $n = 1, \ldots, N$. Given the current configuration of labels, groups and clustering variables, and supposing that the current $Z_{ng} = 1$, the full conditional distribution of $\mathbf{Z}_n$ is found through the relation

$$\begin{aligned} & p(Z_{nh} = 1 | G, \boldsymbol{\nu}, \mathbf{X}, \alpha, \pi, \beta) \\ \propto \; & \Gamma\{N_h - 1 + \alpha\} \\ \times \; & \prod_{m \in \nu_{\text{cl}}} \frac{\prod_{c=1}^{C_m} \Gamma\{n_{hmc} - \mathrm{I}(X_{nm} = c) + \beta\}}{\Gamma\{N_h - 1 + C_m\beta\}} \\ \times \; & \Gamma\{N_h + 1 + \alpha\} \prod_{m \in \nu_{\text{cl}}} \\ \times \; & \frac{\prod_{c=1}^{C_m} \Gamma\{n_{hmc} + \mathrm{I}(X_{nm} = c) + \beta\}}{\Gamma\{N_h + 1 + C_m\beta\}}, \end{aligned}$$

for $h \neq g$. Memberships are updated for each observation at each iteration.

3.4.2 Number of classes

The number of classes may be sampled using the approach introduced by Nobile and Fearnside (2007) for Gaussian mixtures. We outline this approach here. Given that there are currently G components, two Metroplis-Hastings moves may be proposed: either the observations assigned to a component are randomly divided into two groups, so that a new component is "ejected" from an existing one, and the number of groups is increased to $G + 1$, or the observations in two distinct groups are merged together, so that a component is "absorbed", and the number of groups decreases to $G - 1$.

A component is added or removed with probability $p_G = 0.5$ or $1 - p_G = 0.5$, except at the endpoints $G \in \{1, G_{\max}\}$, where they are modified appropriately. A component k is chosen at random to "eject" a new component from. To eject the new component, a draw $u \sim \text{Beta}(a, a)$ is made, and each element of the ejecting component k is assigned to the new component $G + 1$ with probability u, otherwise it remains in k. The choice of shape parameter a can have a strong effect on sampler performance, and is most effective when close to empty components are proposed often. This choice is determined by the size of the proposed component, and a suitable value may be obtained by, for example, numerical programming. We refer to Appendix A3 in Nobile and Fearnside (2007) for further details.

If the proposed components' quantities are denoted by a $\tilde{}$, then the acceptance probability

of an eject is $\min(1, A)$ where

$$\begin{aligned} A &= \frac{p(G+1, \tilde{\mathbf{Z}}, \boldsymbol{\nu}|\mathbf{X}, \alpha, \pi, \beta)}{p(G, \mathbf{Z}, \boldsymbol{\nu}|\mathbf{X}, \alpha, \pi, \beta)} \times \frac{1 - p_{G+1}}{p_G} \\ &\times \frac{\Gamma\{a\}\Gamma\{a\}}{\Gamma\{2a\}} \times \frac{\Gamma\{2a + N_k\}}{\Gamma\{a + \tilde{N}_k\}\Gamma\{a + \tilde{N}_{G+1}\}}. \end{aligned}$$

If the move is accepted, an additional label swap between the ejected component and another of the components selected at random is carried out. This is to improve the mixing properties of the sampler, as it means that the eject and its reverse absorption move are not as systematic. For the reverse absorption move, one chooses two components k and k' at random, and places all items from k into k', computing the acceptance probability as $\min(1, A^{-1})$. If accepted, the number of components decreases from $G + 1$ to G.

3.4.3 Clustering variables

To sample the clustering variables an element ν_j is chosen randomly from the variable index $\boldsymbol{\nu}$. If $\nu_j = 0$, it is proposed to add variable j to the collection of clustering variables $\boldsymbol{\nu}_{\text{cl}}$. Otherwise, if $\nu_j = 1$ the proposed move is to take variable j out from the clustering variables and put it in $\boldsymbol{\nu}_{\text{n}}$. The acceptance probability is $\min(1, R)$ with

$$\begin{aligned} R &= \frac{p(G, \mathbf{Z}, \tilde{\boldsymbol{\nu}}|\mathbf{X}, \alpha, \pi, \beta)}{p(G, \mathbf{Z}, \boldsymbol{\nu}|\mathbf{X}, \alpha, \pi, \beta)} \\ &= \left(\frac{\Gamma\{C_j\beta\}}{\Gamma\{\beta\}^{C_j}}\right)^{G-1} \prod_{g=1}^{G} \frac{\prod_{c=1}^{C_j} \Gamma\{N_{gjc} + \beta\}}{\Gamma\{N_g + C_m\beta\}} \\ &\times \left(\frac{\prod_{c=1}^{C_j} \Gamma\{N_{jc} + \beta\}}{\Gamma\{N + C_j\beta\}}\right)^{-1} \end{aligned}$$

if $\nu_j = 0$, with the calculation inverted if $\nu_j = 1$.

3.4.4 Sampling scheme summary

Each sweep of the algorithm thus takes the following three steps:

(a) Update the class membership of observations using a Gibbs sampling step.

(b) Propose to add a component with probability $p_G = 0.5$, otherwise propose to absorb (remove) a component.

(c) Choose one variable at random. If it is not included to cluster variables, propose to do so. If currently included as a clustering variable, propose to exclude it.

The iterations used for posterior inference are taken after an initial burn in phase. In the applications described in Section 5, diagnostic tools in the **R** package **coda** (Plummer et al, 2006) werewere applied to the the log posterior to help determine whether a sufficient number of samples have been run for burn in to occur, and whether thinning of the resulting samples is required. Visually, trace plots of the log posterior, number of groups and number of variables included can also be used to assess the effectiveness of the sampler.

4 Post-hoc procedures for inference

The MCMC output from running the algorithm described in Section 3.4 can be post-processed in order to perform inference on the clustering of observations, as well as item probability and weight parameter estimation. The first step is to correct the output samples of class labels for label switching. How this is done is outlined in the following. Note that no label switching is required in order to evaluate the number of groups which underly the data, or which variables are useful for clustering. After this, other post-hoc procedures for inference are discussed.

4.1 Label switching

Label switching can occur in the algorithm of Section 3.4. The reason is that

$$p(G, \mathbf{Z}, \boldsymbol{\nu}|\mathbf{X}, \alpha, \pi, \beta) = p(G, \mathbf{Z}_{\cdot\delta}, \boldsymbol{\nu}|\mathbf{X}_n, \alpha, \pi, \beta)$$

where $\mathbf{Z}_{\cdot\delta}$ denotes the indicator matrix obtained by applying any permutation δ of $1, \ldots, G$ to the columns in $\mathbf{Z}$. This invariance to permutations of the labels makes posterior inference of the clusterings fruitless unless some post-processing procedure is employed to try to "undo" the label switching first (Stephens, 2000; Celeux et al, 2000; Marin et al, 2005). Here, the procedure used is the same as that proposed by Nobile and Fearnside (2007) and discussed in detail in Wyse and Friel (2012). Here we provide a brief outline.

The method re-labels samples by minimising a cost function of the group membership vectors $\mathbf{Z}$. Let $\mathbf{Z}^{(T)}$ denote the value of the group membership indicator matrix $\mathbf{Z}$ stored at iteration T during the sample run. Then a $G \times G$ cost matrix $\mathbf{C}$ can be created with entries

$$C_{gh} = \sum_{t=1}^{T-1} \sum_{n=1}^{N} (1 - Z_{ng}^{(t)}) Z_{nh}^{(T)}.$$

A cost function for $\mathbf{C}$ can then be constructed which is minimised by the permutation δ of $\mathbf{Z}_n^{(T)}$, $n = 1, \ldots, N$, which minimises the trace of $\mathbf{C}$. This is found using the square assignment algorithm of Carpaneto and Toth (1980).

4.2 Post-hoc parameter estimation

Although we integrate out the mixture weight and item probability parameters $\boldsymbol{\tau}$ and $\boldsymbol{\theta}$ from the model, it is possible to estimate the *a posteriori* expectations and standard deviations, conditional on a given number of groups, from a post-hoc calculation, by making use of the following formulae:

$$\begin{aligned} \mathbb{E}[A] &= \mathbb{E}[\mathbb{E}[A|B]] \\ \mathbb{V}\text{ar}[A] &= \mathbb{E}[\mathbb{V}\text{ar}[A|B]] + \mathbb{V}\text{ar}[\mathbb{E}[A|B]], \end{aligned}$$

for any random variables A and B.

Define the following summary statistics, for $t \in 1, \ldots, T$ iterations:

$$\begin{aligned} N_g^{(t)} &:= \sum_{n=1}^{N} Z_{ng}^{(t)} \\ S_{gmc}^{(t)} &:= \sum_{n=1}^{N} Z_{ng}^{(t)} \mathrm{I}(X_{nm} = c). \end{aligned}$$

Then we can estimate the expected values

$$\begin{aligned} \mathbb{E}[\theta_{gmc}|\mathbf{X}] &= \mathbb{E}[\mathbb{E}[\theta_{gmc}|\mathbf{X}, \mathbf{Z}]] \\ &\approx \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}[\theta_{gmc}|\mathbf{X}, \mathbf{Z}^{(t)}], \end{aligned}$$

and since $p(\boldsymbol{\theta}_{gm}|\mathbf{X}, \mathbf{Z}^{(t)}, \beta)$ follows a Dirichlet distribution,

$$\mathbb{E}[\theta_{gmc}|X, \beta] \approx \frac{1}{T} \sum_{t=1}^{T} \frac{S_{gmc}^{(t)} + \beta}{N_g^{(t)} + C_m \beta}.$$

It is easy to show based on a similar calculation that

$$\mathbb{E}[\tau_g|\alpha] \approx \frac{1}{T} \sum_{t=1}^{T} \frac{N_g^{(t)} + \alpha}{N + G\alpha}.$$

To calculate the variances:

$$\begin{aligned} & \mathbb{V}\mathrm{ar}[\theta_{gmc}|\mathbf{X}, \beta] \\ = \; & \mathbb{E}[\mathbb{V}\mathrm{ar}[\theta_{gmc}|\mathbf{X}, \mathbf{Z}, \beta]] + \mathbb{V}\mathrm{ar}[\mathbb{E}[\theta_{gmc}|\mathbf{X}, \mathbf{Z}, \beta]] \\ \approx \; & \frac{1}{T} \sum_{t=1}^{T} \frac{(S_{gmc}^{(t)} + \beta)(N_g^{(t)} + (C_m - 1)\beta - S_{gmc}^{(t)})}{(N_g^{(t)} + C_m\beta)^2 (N_g^{(t)} + C_m\beta + 1)} \\ + \; & \frac{1}{T} \sum_{t=1}^{T} \left(\frac{S_{gmc}^{(t)} + \beta}{N_g^{(t)} + C_m\beta} - \frac{1}{T} \sum_{t=1}^{T} \frac{S_{gmc}^{(t)} + \beta}{N_g^{(t)} + C_m\beta} \right)^2, \end{aligned}$$

where we have again made use of the fact that
$p(\boldsymbol{\theta}_{gm}|\mathbf{X}, \mathbf{Z}^{(t)}, \beta)$ follows a Dirichlet distribution to calculate the variance $\mathbb{V}\mathrm{ar}[\theta_{gmc}|\mathbf{X}, \mathbf{Z}^{(t)}, \beta]$.

Similarly,

$$\begin{aligned} \mathbb{V}\mathrm{ar}[\tau_g] \approx \; & \frac{1}{T} \sum_{t=1}^{T} \frac{(N_g^{(t)} + \alpha)(N - N_g^{(t)} + (G-1)\alpha)}{(N + G\alpha)^2 (N + G\alpha + 1)} \\ + \; & \frac{1}{T} \sum_{t=1}^{T} \left(\frac{N_g^{(t)} + \alpha}{N + G\alpha} - \frac{1}{T} \sum_{t=1}^{T} \frac{N_g^{(t)} + \alpha}{N + G\alpha} \right)^2. \end{aligned}$$

4.3 Model choice using Bayes factors

Once a satisfactory number of MCMC samples have been obtained, it still remains to choose an optimal model for the data. Broadly, there are two choices to be made: the optimal number of groups which underly the data, and identify which variables are relevant to the clustering. These choices can be made separately by considering the marginal distributions of the parameters G and $\boldsymbol{\nu}$ respectively.

While it is straighforward to identify an optimal model using the algorithms proposed in Section 3.4 (this will simply be the model occuring most frequently during the sampling scheme), ideally it would also be possible to compare the different levels of evidence for competing models in a manner consistent with the existing literature. For example, it is important to differentiate between the case where there is a large body of evidence in favour of a particular model in comparison with others, as opposed to the case where several competing models may be equally suitable. We motivate our model choices based on the use of Bayes factors (Kass and Raftery, 1995). This is a similar approach to Dean and Raftery (2010), except that the Bayes factors are calculated directly based on MCMC estimates of the marginal probabilities for a particular model, rather than using the BIC criterion as an approximation.

When comparing a model $\mathcal{M}_i$, against a model $\mathcal{M}_j$, for some data $\mathbf{X}$ the posterior odds in favor of model $\mathcal{M}_i$ are given to be

$$\frac{p(\mathcal{M}_i|\mathbf{X})}{p(\mathcal{M}_j|\mathbf{X})} = \frac{p(\mathbf{X}|\mathcal{M}_i)}{p(\mathbf{X}|\mathcal{M}_j)} \times \frac{p(\mathcal{M}_i)}{p(\mathcal{M}_j)},$$

or in words, that the posterior odds are equal to the Bayes factor times the prior odds. It follows immediately that the Bayes factor $\mathcal{B}_{i,j}$, for evaluating the evidence in favour of $\mathcal{M}_i$ versus $\mathcal{M}_j$, can be calculated as:

$$\mathcal{B}_{i,j} = \frac{p(\mathcal{M}_i|\mathbf{X})}{p(\mathcal{M}_j|\mathbf{X})} \times \frac{p(\mathcal{M}_j)}{p(\mathcal{M}_i)},$$

or, in words, the posterior odds times the inverted prior odds.

We first calculate Bayes factors for group choice. Let g be the true number of groups in our model. We have that, for model $\mathcal{M}_k := \{\text{Model with } g = k\}$,

$$\mathbb{P}(g = k) := P_k = \frac{p(\mathbf{X}, \mathcal{M}_k)}{\sum_{j=1}^{G_{\max}} p(\mathbf{X}, \mathcal{M}_j)}.$$

Then posterior odds between models $\mathcal{M}_k$ and $\mathcal{M}_{k-1}$ can then be calculated:

$$\begin{aligned} \frac{P_k}{P_{k-1}} &= \frac{p(\mathbf{X}, \mathcal{M}_k)}{p(\mathbf{X}, \mathcal{M}_{k-1})} \\ &= \frac{p(\mathcal{M}_k|\mathbf{X})}{p(\mathcal{M}_{k-1}|\mathbf{X})}. \end{aligned}$$

The prior odds between models $\mathcal{M}_k$ and $\mathcal{M}_{k-1}$ are also straightforward to find. As noted previously in Section 3, based on an upper limit $G_{\max}$, the prior distribution for $\mathcal{M}_k$ is a truncated Pois(1) distribution. Then $p(\mathcal{M}_k) = (1/k!) \,/\, \left(\sum_{j=1}^{G_{\max}} 1/j!\right)$.

The prior odds for models $\mathcal{M}_k$ and $\mathcal{M}_{k-1}$ are then

$$\frac{p(\mathcal{M}_k)}{p(\mathcal{M}_{k-1})} = \frac{1/k!}{1/(k-1)!} = \frac{(k-1)!}{k!} = \frac{1}{k}.$$

It follows from these two calculations that

$$\mathcal{B}_{k,k-1} = \frac{P_k}{P_{k-1}} \times k.$$

We follow Kass and Raftery (1995), in that values of $\mathcal{B}_{k,k-1}$ between 1 and 3 are viewed as barely worth mentioning, between 3 and 10 as substantial, and greater than 10 as strong. Values greater than 100 are viewed as decisive.

When determining whether to include or exclude a variable, calculation of the Bayes factor is even more straightforward. Since the *a priori* probability of inclusion for any variable has been set as $\pi = 0.5$, it follows that the Bayes factor for variable inclusion is equal to the posterior odds.

Defining $\mathcal{M}_d := \{\text{Model including variable } d\}$, and $\mathcal{M}_{\bar{d}} := \{\text{Model excluding variable } d\}$, then

$$\mathbb{P}(\text{variable } d \text{ included}) := P_d = \frac{p(\mathbf{X}, \mathcal{M}_d)}{p(\mathbf{X}, \mathcal{M}_d) + p(\mathbf{X}, \mathcal{M}_{\bar{d}})}.$$

It follows that

$$\mathcal{B}_{d,\bar{d}} = \frac{P_d}{P_{\bar{d}}}.$$

Since the Bayes factor in this case corresponds directly to its posterior probability, it is particularly straightforward to assess; in particular, the Bayes factor for a variable's inclusion will be greater than 1, 3 and 10 when the variable is included for more than 50%, 75% and 90.9% of the samples respectively. When determining whether or not to include a variable, Dean and Raftery (2010) choose a threshold equivalent to $\mathcal{B}_{d,\bar{d}} > 1$. This corresponds to including any variable which is included more often than it is excluded during a sampling run.

5 Data Applications

In this section the sampler is applied to the datasets described in Section 2.2. In all cases the sampler was first run on the data for 5,000 iterations after 1,000 iterations burn in. A better individually tuned sampler was then run on each dataset, based on these initial runs. Multiple runs of the sampler were also performed on each dataset to ensure consistency of results.

The algorithms described in Sections 3 and 4 were implemented in C and applied to the datasets in the R environment (R Core Team, 2013). Total run time varied depending on the size and particularly the number of variables of the dataset in question. For example, the total time taken to fit a model to the Alzheimer dataset, including post hoc parameter estimation, was about three and a half minutes. Fitting a finely tuned collapsed sampler

to the larger Teaching dataset, without performing post hoc parameter estimation, took about 25 minutes, or about seven and a half times as long. Nevertheless, this should still be viewed as being reasonably efficient, as model and variable selection, as well as clustering, is simultaneously taking place.

In order to visually summarise the sampler's behaviour and simultaneously identify which variables and number of groups are of most interest we construct an $M \times G_{\max}$ coincidence matrix $\boldsymbol{\mathcal{C}}$. Each entry indicates the amount of time which the sampler spent for a certain group number and including a certain variable, and is calculated as follows.

Let $\boldsymbol{\mathcal{G}}$ be a $G_{\max} \times T$ indicator matrix, where T denotes the total number of iterations which the sampler runs for. We define an entry of $\boldsymbol{\mathcal{G}}$ to be:

$$\mathcal{G}_{kt} = \begin{cases} 1 & \text{if the chain has } k \text{ groups at iteration } t; \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, we use $\boldsymbol{\mathcal{V}}$ to denote an $M \times T$ indicator matrix, where

$$\mathcal{V}_{mt} = \begin{cases} 1 & \text{if } \nu_m^{(t)} = 1; \\ 0 & \text{otherwise.} \end{cases}$$

Then each entry of $\boldsymbol{\mathcal{C}}$ is given by

$$\mathcal{C}_{km} = \frac{\sum_{t=1}^{T} \mathcal{V}_{mt}\mathcal{G}_{kt}}{\sum_{t=1}^{T} \mathcal{G}_{kt}}$$

where we have normalised the entry $\mathcal{C}_{km}$ so that it denotes the proportion, rather than the total amount of time the sampler spent in a particular model space. While in theory the matrix summarises the behaviour of the sampler for the entire model space, in practice some regions will not be visited by the sampler, with the corresponding entries being omitted in what follows.

5.1 Dean and Raftery data

A sampler was run on the data for 50,000 iterations after 1,000 iterations burn in, with 1 in 10 samples retained by successive subsampling. The coincidence matrix is shown in Figure 1(a). This clearly identifies a two group model with variables 1-4 as being optimal for clustering. In particular, variables 5-13 are included less than half the time, whereas variables 1-4 are included with high probability. The choice of two groups is decisively confirmed in Table 2, with a Bayes factor of 18.9 in favour of two groups compared to one, and small Bayes factors for all other models. The sampler correctly classifies 381 of the 500 observations, only seven less than would be found using the true parameter values. These results are comparable to those found by Dean and Raftery (2010).

5.2 Alzheimer Data

Initially, the sampler was run on the Alzheimer data for 20,000 iterations after 1,000 iterations burn in. While a visual inspection of the log posterior suggested that good mixing was

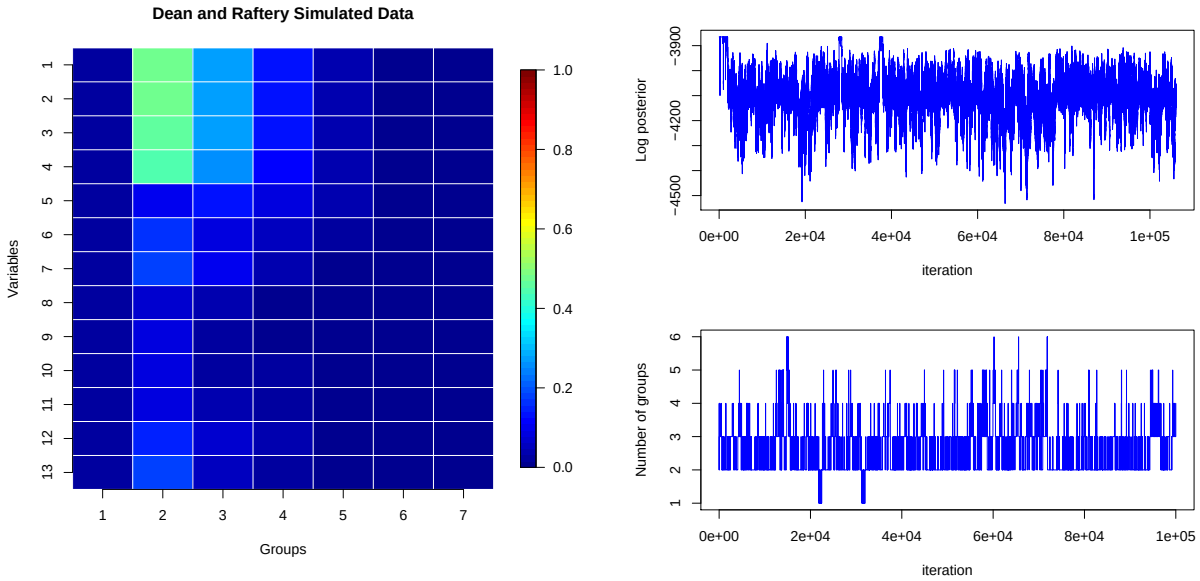

Figure 1: Summary of the spaces most commonly visited by the sampler over a run on simulated data. Right: a trace plot of the log-posterior and the number of groups from a run of the sampler on simulated data.

Table 2: Bayes factors for group choice for the Dean and Raftery data. A two group model is correctly identified as being optimal.

$\mathcal{BF}_{1,2}$	$\mathcal{BF}_{2,3}$	$\mathcal{BF}_{3,4}$	$\mathcal{BF}_{4,5}$	$\mathcal{BF}_{5,6}$	$\mathcal{BF}_{6,7}$
18.88	1.77	1.91	1.46	0.94	2.26

occurring, the log posterior samples was found to have high autocorrelation, and diagnostic tools suggested that a longer sampling run was required. The sampler was then run for 100,000 iterations, and thinned by successively selecting at random one of twenty samples. While this substantially reduced the amount of autocorrelation between samples, the results of the clustering remained relatively unaffected. Diagnostic plots of the sampling run are shown in Figure 2.

The coincidence matrix of the sampling run is shown in the rightmost plot of Figure 2. The sampler excluded variable 1, Hallucination, a majority of the time, and identified 2 or 3 groups as optimal for the data, spending a majority of time in the 2 group space. The Bayes factor for the evidence in favour of a 3 instead of group model 2 is $\mathcal{BF}_{2,3} = 1.43$, suggesting that there is little evidence in favour of the larger group choice. The Bayes factors for other group choices are smaller again. The Bayes factors for group choice are given in Table 3. The Bayes factors for variable selection are given in the supplementary material.

Table 3: Bayes factors for group choice for the Alzheimer data.

$\mathcal{BF}_{2,3}$	$\mathcal{BF}_{3,4}$	$\mathcal{BF}_{4,5}$	$\mathcal{BF}_{5,6}$
1.43	0.83	0.77	0.12

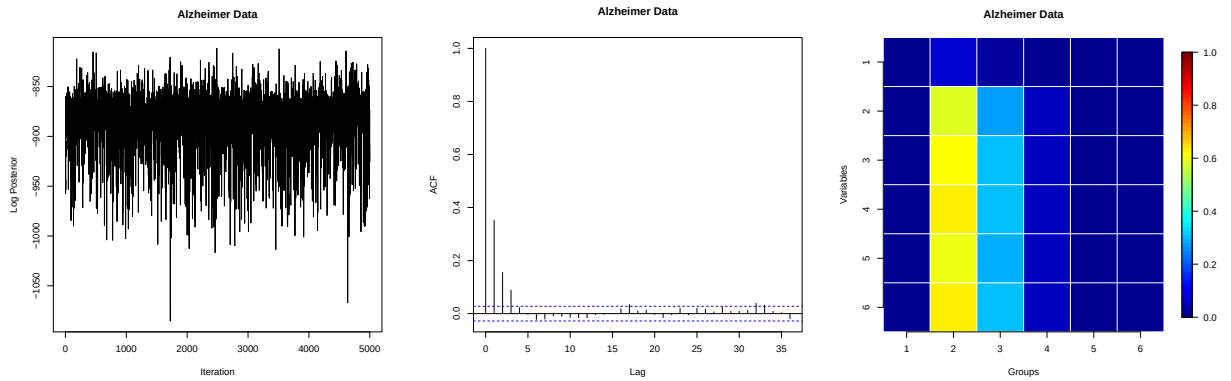

Figure 2: Plots of the collapsed sampler applied to the Alzheimer dataset. Left and centre are diagnostic plots. On the right, a coincidence matrix for the sampler applied to Alzheimer dataset. The model appears to identify 2 groups as optimal, and frequently excluded variable 1 (Hallucination) during the run.

We also compare parameter maximum *a posteriori* and posterior standard deviation estimates for the item probability parameters of a 2 group model fitted to this dataset using the collapsed sampler to those obtained using a full Gibbs sampler in the **BayesLCA** package. See Garrett and Zeger (2000) and Walsh (2006) for descriptions of a full model Gibbs sampler method for LCA. Estimates from the full Gibbs sampler were obtained from 50,000 iterations after 1,000 iterations burn-in and, with every tenth sample retained. Estimates from the two methods are nearly identical. These are given in Table 4.

5.3 Teaching Styles Data

A collapsed sampler was run on the Teaching Styles data for 200,000 iterations after 5,000 iterations burn in, with 1 in 20 samples retained by successive subsampling. While a certain amount of lag in the auto correlation persists in the log posterior samples, diagnostic tools suggest that the sampler behaves satisfactorily on the data. Diagnostic plots of the sampler are shown in Figure 3.

The coincidence matrix for the sampler is also shown as the rightmost plot in Figure 3. While several variables are decisively dropped by the sampler, a certain amount of uncertainty exists when it comes to selecting an optimal number of groups. Bayes factors for the model are given in Table 5. The largest Bayes factor indicates a large amount of evidence in favour of a 7 group model for the data in favour of 6 group model, $\mathcal{BF}_{6,7} = 13.16$. There is also some evidence for an 8 group model, with $\mathcal{BF}_{7,8} = 5.4$. Both the 7 and 8 group models cluster some observations into quite small groups; one group in the 7 group model maps only 10 observations to one group, while the two smallest groups in the 8 group model contain only 2 and 15 observations respectively.

A cross-classification table comparison of the 6 and 7 group models is shown in Table 6. This shows that the main difference between the two models from a clustering perspective is the introduction of an extra group - Group 6 in the 7 group model - which contains

Table 4: MAP item probability estimates (posterior standard deviations in parentheses) for the Alzheimer data obtained from post-hoc estimates from the collapsed Gibbs sampler and a from full model Gibbs sampler run. The parameter mean and standard deviation estimates are nearly identical.

Collapsed Gibbs sampler post-hoc estimates

	Hallucination	Activity	Aggression	Agitation	Diurnal	Affective
Group 1	0.08 (0.03)	0.54 (0.06)	0.10 (0.04)	0.14 (0.06)	0.13 (0.05)	0.59 (0.08)
Group 2	0.10 (0.04)	0.80 (0.06)	0.40 (0.08)	0.64 (0.12)	0.39 (0.07)	0.94 (0.04)

Full model Gibbs sampler estimates

	Hallucination	Activity	Aggression	Agitation	Diurnal	Affective
Group 1	0.08 (0.03)	0.54 (0.06)	0.11 (0.05)	0.14 (0.06)	0.14 (0.05)	0.59 (0.08)
Group 2	0.10 (0.04)	0.79 (0.07)	0.39 (0.08)	0.64 (0.12)	0.38 (0.07)	0.93 (0.07)

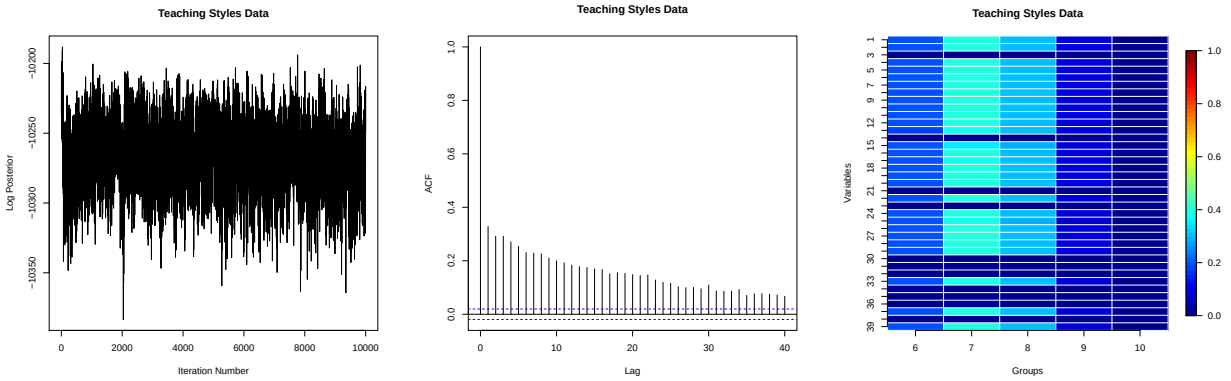

Figure 3: Diagnostic plots of the collapsed sampler applied to the Teaching Styles dataset, and on the right, a coincidence matrix for the sampler.

observations clustered to Groups 1 and 3 with high uncertainty in the 6 group model. This suggests that the additional groups in the model may improve model fit without necessarily introducing additional distinct clusters to the data.

Table 5: Bayes factors for number of groups in Teaching Styles data.

$\mathcal{BF}_{6,7}$	$\mathcal{BF}_{7,8}$	$\mathcal{BF}_{8,9}$	$\mathcal{BF}_{9,10}$	$\mathcal{BF}_{10,11}$
13.16	5.40	3.84	2.29	1.32

Figure 4 compares the dropped variables to those retained by the sampler. This shows heatmap plots of the item probability parameters for the 7 group model, firstly for the included and then the excluded variables. (Note that due to space restrictions, only every second variable index is included for the left plot.) Parameter estimates are highly similar for the excluded variables, while the behaviour of the parameters is much more varied for the included variables. While the estimates of parameters in Group 3 appear different from

Table 6: Comparison of mapped clusterings for the 6 and 7 group models fitted to the Teaching Styles dataset. For convenience some rows in the table have been re-arranged to illustrate the agreement in the clusterings.

	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
Group 1	40	0	3	0	0	0
Group 2	0	59	0	0	0	0
Group 3	0	0	105	0	0	0
Group 4	1	0	6	141	0	0
Group 5	0	0	3	0	51	0
Group 6	6	0	4	0	0	0
Group 7	0	0	1	0	0	47

other groups for the excluded parameters, this is the smallest group in the clustering, and as a result standard error estimates for this group's item probability parameters are extremely high, ranging between 35% and 43%.

5.4 Physiotherapy Data

A collapsed sampler was run on the Physiotherapy data for 20,000 iterations after 5,000 iterations burn in, with every second iteration retained. Auto correlation of the sampler was deemed satisfactory after these measures were taken. Diagnostic plots of the sampler are provided in the supplementary material.

When applied to the full dataset, the sampler settles on 7 groups. (While a certain amount of evidence exists for an 8 group model, when compared, the two groupings have 98% agreement.) This is shown in Table 7. These 7 groups correspond well with the three pain types. A confusion matrix comparing the grouped observations with their reported pain types is shown in Table 8. Overall there is high agreement between the two sets of groupings, with a Rand index of about 92%. Note that some groups have extremely small mapped membership levels. Almost all variables proved useful in this clustering, with only one variable being dropped by the sampler. This is perhaps unsurprising, since the variables were identified for the study by expert recommendation.

Table 7: Bayes factors for number of groups in Physiotherapy data.

$\mathcal{BF}_{6,7}$	$\mathcal{BF}_{7,8}$	$\mathcal{BF}_{8,9}$	$\mathcal{BF}_{9,10}$
613.66	5.14	0.97	0.72

A further study was then made to see which variables proved useful for clustering within each pain type. Samplers of length 75,000, 20,000 and 75,000 were run after 5,000 iterations of burn in for observations reporting central neuropathic, nociceptive and peripheral neuropathic pain types respectively. For the samplers run on the central neuropathic and peripheral neuropathic datasets, the chains were thinned in a highly conservative manner, with 1 in 15 samples retained by successive subsampling. Every second iteration was retained for the nociceptive data.

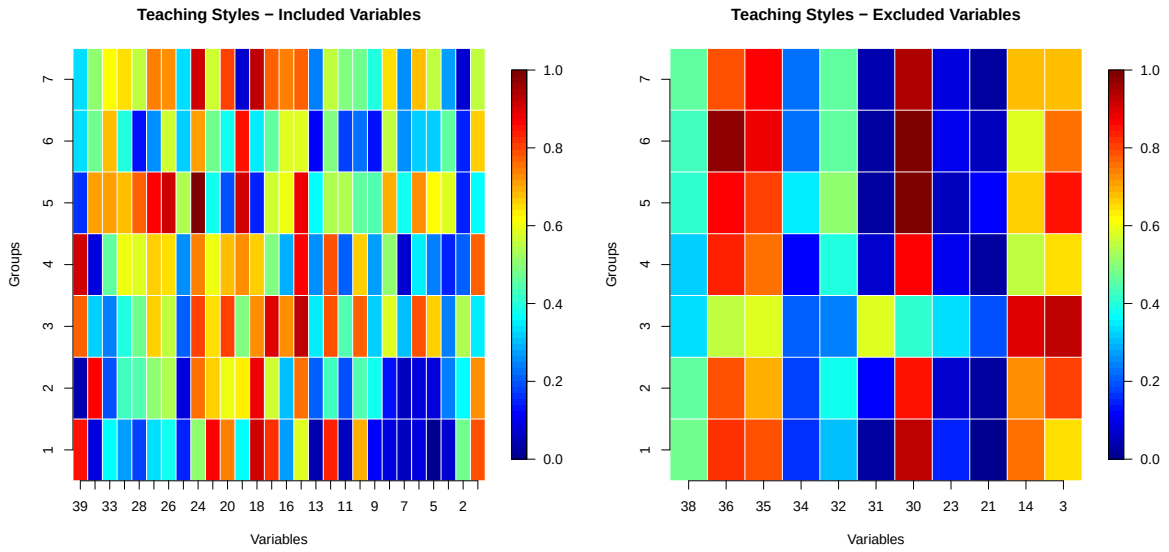
Teaching Styles – Included Variables
Groups
Variables
39 33 28 26 24 20 18 16 13 11 9 7 5 2
1 2 3 4 5 6 7
1.0 0.8 0.6 0.4 0.2 0.0
Teaching Styles – Excluded Variables
Groups
Variables
38 36 35 34 32 31 30 23 21 14 3
1 2 3 4 5 6 7
1.0 0.8 0.6 0.4 0.2 0.0

Figure 4: Heatmap plots of the item probability parameters for 7 group model fitted to the Teaching Styles data, for the included and excluded variables. Parameter estimates are highly similar for the excluded variables, while varied behaviour is apparent for included variables.

In all cases, a high number of variables were decisively dropped by the samplers: 9, 15, and 18 variables were dropped for the respective datasets, with strong evidence that 21, 17 and 12 of the variables were informative to the clustering. No variables were excluded by all three models. The samplers identify 5, 4 and 3 group models as suitable for the data subsets, although again, several of the groups are quite small, particularly for the model fitted to the nociceptive data.

Comparing the clusterings of the models fitted to the data subsets to the clustering of the relevant observations from the model fitted to the full dataset, shows a high level of agreement, with Rand indices of 82%, 90% and 93% respectively. In particular, it is worth noting that for the models fitted to the central neuropathic and Nociceptive subsets, the same number of groups are chosen as observations for the respective pain type are clustered into by the model fitted to the full dataset. For the peripheral neuropathic subset, only two observations with this pain type are not placed into one of three groups, the number of groups chosen by the model fitted to the data subset.

The main difference between the clusterings is that the groups in models fitted to the data subsets are of a more equal size in comparison with the relevant subsets obtained from clustering the full dataset; this makes sense, as excluding the variables which are not meaningful for the subset in question make it easier to distinguish between different types of behaviour. A comparison of the clusterings for the nociceptive dataset is shown in Table 9.

Table 8: Confusion matrix comparing clustered observations to reported pain types in the Physiotherapy data. (CN = Central Neuropathic, N=Nociceptive, and PN = Peripheral Neuropathic.) For convenience some rows in the table have been re-arranged to illustrate the agreement in the clusterings.

	Pain Type		
	CN	N	PN
Group 1	3	0	1
Group 5	52	0	0
Group 7	30	3	0
Group 3	6	96	1
Group 6	0	120	1
Group 2	1	16	79
Group 4	3	0	13

Table 9: Confusion matrix comparing clustered observations based on the model fitted to the nociceptive subset to the clustering of the relevant observations obtained from the 7 group model applied to the full Physiotherapy dataset.

		Full dataset			
		Group 6	Group 3	Group 2	Group 7
Nociceptive subset	Group 2	107	2	0	0
	Group 4	0	81	0	1
	Group 1	13	4	15	0
	Group 3	0	9	1	2

6 Discussion

In this paper we have introduced a model to perform LCA, including group and variable selection, in a Bayesian setting. Collapsing of some of the model parameters leads to a more efficient algorithm, where the clustering of observations is emphasised, and if required, accurate estimates for the marginalised parameters can be obtained from a post hoc procedure. Directly sampling the number of groups and the inclusion of variables allows for a suitable model to be chosen in a highly principled manner.

In the latter two applications described in Section 5, some groups in the selected model contained only a small number of observations. These observations might potentially be viewed as outliers, or as being poorly described by the larger groups found by the model, rather than as distinct clusters. This is a particularly pragmatic approach to take when clustering in a model based setting, where the resultant high posterior standard deviations of model parameters for these small groups make interpretation of the group behaviour essentially meaningless.

We note that the outlined model for variable selection is somewhat limited by the conditional independence assumption of the LCA model. In words, when a variable is proposed to be included or excluded by the sampling scheme, we are asking the question "does the proposed variable contain information about the clustering?", rather than "does the proposed variable contain *additional* information about the clustering?" Thus, while non-informative

variables are removed satisfactorily, for example, two informative but highly correlated variables would both be included with high probability, when perhaps the inclusion of just one would provide clustering results of similar quality. Raftery and Dean (2006) propose such a model for variable selection when clustering continuous data; the covariance parameter of a multivariate normal distribution being a natural way to model the conditional dependence between variables for this type of data.

Latent trait analysis models (see Bartholomew and Knott, 1999, for example) allow dependency between multivariate categorical data, and Gollini and Murphy (2013) have recently proposed a mixture of latent trait analyzers model which does not make the conditional independence assumption of LCA. Potentially this model could prove useful for variable selection in a clustering setting. On the other hand, implementing the model for this purpose may be more difficult than for LCA, as the likelihood is not available in closed form.

Finally, in Section 5.4 we applied the collapsed sampler to subsets of the dataset, based on additionally labelled information. In a general clustering setting, it may be of interest to determine which variables are most helpful for identifying particular groups in the data. This is a sort of unsupervised discriminant analysis which may be of future interest to further reduce the dimensionality of a dataset.

References

Aitkin M, Anderson D, Hinde J (1981) Statistical modelling of data on teaching styles. Journal of the Royal Statistical Society Series A (General) 144(4):pp. 419–461

Akaike H (1973) Information theory and an extension of the maximum likelihood principle. In: Second International Symposium on Information Theory, Akadémiai Kiadó, pp 267–281

Bartholomew DJ, Knott M (1999) Latent Variable Models and Factor Analysis, 2nd edn. Kendall's Library of Statistics, Hodder Arnold

Bennet N (1976) Teaching Styles and Pupil Progress. Open Books, London

Carpaneto G, Toth P (1980) Algorithm 548: Solution of the assignment problem [H]. ACM Transactions on Mathematical Software 6(1):104–111

Celeux G, Hurn M, Robert CP (2000) Computational and inferential difficulties with mixture posterior distributions. Journal of the American Statistical Association 95(451):pp. 957–970, URL `http://www.jstor.org/stable/2669477`

Celeux G, Forbes F, Robert CP, Titterington D (2006) Deviance information criteria for missing data models. Bayesian Analysis 1(4):pp. 651–673

Dean N, Raftery AE (2010) Latent class analysis variable selection. The Annals of the Institute of Statistical Mathematics 62(1):11–35

Dellaportas P, Papageorgiou I (2006) Multivariate mixtures of normals with unknown number of components. Statistics and Computing 16(1):57–68

Dempster AP, Laird NM, Rubin DB (1977) Maximum Likelihood from Incomplete Data via the EM Algorithm. Journal of the Royal Statistical Society B 39(1):1–38

Fraley C, Raftery A (2007) Model-based methods of classification: Using the software in chemometrics. Journal of Statistical Software 18(6):1–13

Garrett ES, Zeger SL (2000) Latent class model diagnosis. Biometrics 56(4):pp. 1055–1067

Geman S, Geman D (1984) Stochastic Relaxation, Gibbs Distributions and the Bayesian Restoration of Images. IEEE Transactions on Pattern Analysis and Machine Intelligence 6(6):721–741

Gollini I, Murphy T (2013) Mixture of latent trait analyzers for model-based clustering of categorical data. Statistics and Computing (to appear)

Goodman LA (1974) Exploratory latent structure analysis using both identifiable and unidentifiable models. Biometrika 61(2):215–231

Green P (1995) Reversible jump markov chain monte carlo computation and bayesian model determination. Biometrika 82(4):711–732

Kass RE, Raftery AE (1995) Bayes Factors. Journal of the American Statistical Association 90(430):773–795

Marin JM, Mengersen K, Robert CP (2005) Bayesian modelling and inference on mixtures of distributions. In: Dey D, Rao C (eds) Bayesian Thinking: Modeling and Computation, Handbook of Statistics, vol 25, 1st edn, North Holland, Amsterdam, chap 16, pp 459–507

McDaid AF, Murphy TB, Friel N, Hurley N (2013) Improved Bayesian inference for the stochastic block model with application to large networks. Computational Statistics & Data Analysis 60(0):12 – 31

McLachlan G, Peel D (2002) Finite Mixture Models. John Wiley & Sons

Moran M, Walsh C, Lynch A, Coen RF, Coakley D, Lawlor BA (2004) Syndromes of behavioural and psychological symptoms in mild Alzheimer's disease. International Journal of Geriatric Psychiatry 19(4):359–364

Nobile A, Fearnside A (2007) Bayesian finite mixtures with an unknown number of components: The allocation sampler. Statistics and Computing 17:147–162

Plummer M, Best N, Cowles K, Vines K (2006) CODA: Convergence diagnosis and output analysis for MCMC. R News 6(1):7–11

R Core Team (2013) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, URL `http://www.R-project.org/`

Raftery AE, Dean N (2006) Variable selection for model-based clustering. Journal of the American Statistical Association 101(473):168–178, DOI 10.1198/016214506000000113

Raftery AE, Newton MA, Satagopan JM, Krivitsky PN (2007) Estimating the integrated likelihood via posterior simulation using the harmonic mean identity (with discussion). In: Bernardo J, Bayarri M, Berger J, Dawid A, Heckerman D, Smith A, West M (eds) Bayesian Statistics 8, Oxford University Press, Oxford, pp 1–45

Richardson S, Green PJ (1997) On bayesian analysis of mixtures with an unknown number of components (with discussion). Journal of the Royal Statistical Society: Series B (Statistical Methodology) 59(4):731–792

Rousseau J, Mengersen K (2011) Asymptotic behaviour of the posterior distribution in overfitted mixture models. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 73(5):689–710

Schwarz G (1978) Estimating the Dimension of a Model. The Annals of Statistics 6(2):461–464, DOI 10.2307/2958889

Smart KM, Blake C, Staines A, Doody C (2011) The discriminative validity of "nociceptive", "peripheral neuropathic", and "central sensitization" as mechanisms-based classifications of musculoskeletal pain. Clinical Journal of Pain 27:655 – 663

Spiegelhalter DJ, Best NG, Carlin BP, van der Linde A (2002) Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society B 64(4):583–639

Stephens M (2000) Dealing with Label Switching in Mixture Models. Journal of the Royal Statistical Society B 62(4):795–809

Tadesse MG, Sha N, Vannucci M (2005) Bayesian Variable Selection in Clustering High-Dimensional Data. Journal of American Statistical Association 100(470):602–617

Walsh C (2006) Latent class analysis identification of syndromes in Alzheimer's disease: A Bayesian approach. Metodološki Zvezki - Advances in Methodology and Statistics 3(1):147 – 162

White A, Murphy B (2013) BayesLCA: Bayesian Latent Class Analysis. URL `http://CRAN.R-project.org/package=BayesLCA`, r package version 1.3

Wyse J, Friel N (2012) Block clustering with collapsed latent block models. Statistics and Computing 22(1):415–428